

Phagocytes: A Holistic Defense and Protection Against Active P2P Worms

Ruichuan Chen[†], Eng Keong Lua[‡], Jon Crowcroft[§], Liyong Tang[†], Zhong Chen[†]

[†]Peking University, China; [‡]Carnegie Mellon University, USA; [§]University of Cambridge, UK

Abstract

Active Peer-to-Peer (P2P) worms present serious threats to the global Internet by exploiting popular P2P applications to perform rapid topological self-propagation. Active P2P worms pose more deadly threats than normal scanning worms because they do not exhibit easily detectable anomalies, thus many existing defenses are no longer effective.

We propose an immunity system with Phagocytes — a small subset of elected P2P hosts that are immune with high probability and specialized in finding and “eating” worms in the P2P overlay. The Phagocytes will monitor their managed P2P hosts’ connection patterns and traffic volume in an attempt to detect active P2P worm attacks. Once detected, local isolation, alert propagation and software patching will take place for containment. The Phagocytes further provide the access control and filtering mechanisms for communication establishment between the internal P2P overlay and the external hosts. We design a novel adaptive and interaction-based computational puzzle scheme at the Phagocytes to restrain external worms attacking the P2P overlay, without influencing legitimate hosts’ experiences significantly. We implement a prototype system, and evaluate its performance based on realistic massive-scale P2P network traces. The evaluation results illustrate that our Phagocytes are capable of achieving a total defense against active P2P worms.

1 Introduction

The ability to gain control of a huge amount of Internet hosts could be easily achieved by the exploitation of worms which self-propagate through popular Internet applications and services. Internet worms have already proven their capability of inflicting massive-scale disruption and damage to the Internet infrastructure. These worms employ normal *scanning* as a strategy to find potential vulnerable targets, i.e., they randomly select victims from the IP address space. So far, there have been many existing schemes that are effective in detecting such scanning worms [14], e.g., by capturing the scanning events [18] or by passively detecting

abnormal network traffic activities [11].

In recent years, Peer-to-Peer (P2P) overlay applications have experienced an explosive growth, and now dominate large fractions of both the Internet users and traffic volume [12]; thus, a new type of worms that leverage the popular P2P overlay applications, called *P2P worms*, pose a very serious threat to the Internet [2]. Generally, the P2P worms can be grouped into two categories: *passive* P2P worms and *active* P2P worms. The passive P2P worm attack is generally launched either by copying such worms into a few P2P hosts’ shared folders with attractive names, or by participating into the overlay and responding to queries with the index information of worms. Unable to identify the worm content, normal P2P hosts download these worms unsuspectingly into their own shared folders, from which others may download later without being aware of the threat, thus passively contributing to the worm propagation. The passive P2P worm attack could be mitigated by current patching systems [28] and reputation models [25]. In this paper, we focus on another serious P2P worm: active P2P worm.

The active P2P worms could utilize the P2P overlay applications to retrieve the information of a few vulnerable P2P hosts and then infect these hosts, or as an alternative, these worms are directly released in a hit list of P2P hosts to bootstrap the worm infection. Since the active P2P worms have the capacity of gaining control of the infected P2P hosts, they could perform rapid *topological self-propagation* by spreading themselves to neighboring hosts, and in turn, spreading throughout the whole network to affect the quality of overlay service and finally cause the overlay service to be unusable. The P2P overlay provides an accurate way for worms to find more vulnerable hosts easily without probing randomly selected IP addresses (i.e., low connection failure rate). Moreover, the worm attack traffic could easily blend into normal P2P traffic, so that the active P2P worms will be more deadly than scanning worms. That is, they do not exhibit easily detectable anomalies in the network traffic as scanning worms do, so many existing defenses against scanning worms are no longer effective [32].

Besides the above internal infection in the P2P overlay, the infected P2P hosts could again mount attacks to external hosts. In similar sense, since the P2P overlay applications

are pervasive on today’s Internet, it is also attractive for malicious external hosts to mount attacks against the P2P overlay applications and then employ them as an ideal platform to perform future massive-scale attacks, e.g., botnet attacks.

In this paper, we aim to develop a *holistic* immunity system to provide the mechanisms of both *internal defense* and *external protection* against active P2P worms. In our system, we elect a small subset of P2P overlay nodes, *Phagocytes*, which are immune with high probability and specialized in finding and “eating” active P2P worms. Each Phagocyte in the P2P overlay is assigned to manage a group of P2P hosts. These Phagocytes monitor their managed P2P hosts’ connection patterns and traffic volume in an attempt to detect active P2P worm attacks. Once detected, the local isolation procedure will cut off the links of all the infected P2P hosts. Afterwards, the responsible Phagocyte performs the contagion-based alert propagation to spread worm alerts to the neighboring Phagocytes, and in turn, to other Phagocytes. Here, we adopt a threshold strategy to limit the impact area and enhance the robustness against the malicious alert propagations generated by infected Phagocytes. Finally, the Phagocytes help acquire the software patches and distribute them to the managed P2P hosts. With the above four modules, i.e., detection, local isolation, alert propagation and software patching, our system is capable of preventing internal active P2P worm attacks from being effectively mounted within the P2P overlay network.

The Phagocytes also provide the access control and filtering mechanisms for the connection establishment between the internal P2P overlay and the external hosts. Firstly, the P2P traffic should be contained within the P2P overlay, and we forbid any P2P traffic to leak from the P2P overlay to external hosts. This is because such P2P traffic is generally considered to be malicious and it is possible that the P2P worms ride on such P2P traffic to spread to the external hosts. Secondly, in order to prevent external worms from attacking the P2P overlay, we hide the P2P hosts’ IP addresses with the help of scalable distributed DNS service, e.g., CoDoNS [20]. An external host who wants to gain access to the P2P overlay has no alternative but to perform an interaction towards the associated Phagocyte to solve an adaptive computational puzzle; then, according to the authenticity of the puzzle solution, the Phagocyte can determine whether to process the request.

We implement a prototype system, and evaluate its performance on a massive-scale testbed with realistic P2P network traces. The evaluation results validate the effectiveness and efficiency of our proposed holistic immunity system against active P2P worms.

Outline. We specify the system architecture in section 2. Sections 3 and 4 elaborate the internal defense and external protection mechanisms, respectively. We then present the experimental design in section 5, and discuss the evaluation




Figure 1: System Architecture

results in section 6. Finally, we give an overview of related work in section 7, and conclude this paper in section 8.

2 System Architecture

Current P2P overlay networks can generally be grouped into two categories [15]: *structured* overlay networks, e.g., Chord [22], whose network topology is tightly controlled based on distributed hash table, and *unstructured* overlay networks, e.g., Gnutella [1], which merely impose loose structure on the topology. In particular, most modern unstructured P2P overlay networks utilize a two-tier structure to improve their scalability: a subset of peers, called *ultra-peers*, construct an unstructured mesh while the other peers, called *leaf-peers*, connect to the ultra-peer tier for participating into the overlay network.

As shown in Figure 1, the network architecture of our system is similar to that of the two-tier unstructured P2P overlay networks. In our system, a set of P2P hosts act as the Phagocytes to perform the functions of defense and protection against active P2P worms. These Phagocytes are elected among the participating P2P hosts in terms of the following metrics: high bandwidth, powerful processing resource, sufficient uptime, and applying the latest patches (interestingly, the experimental result shown in section 6 indicates that we actually do not need to have a large percentage of Phagocytes applying the latest patches). As existing two-tier unstructured overlay networks do, the Phagocyte election is performed periodically; moreover, even if an elected Phagocyte has been infected, our internal defense mechanism (described in section 3) can still isolate and patch the infected Phagocyte immediately. In particular, the population of Phagocytes should be small as compared to the total overlay population, otherwise the scalability and applicability are questionable.

As a result, each elected Phagocyte covers a number of managed P2P hosts, and each managed P2P host will belong to one closest Phagocyte. That is, the Phagocyte acts as the proxy for its managed hosts to participate into the P2P overlay network, and has the control over the managed P2P hosts. Moreover, a Phagocyte further connects to several nearby Phagocytes based on close proximity.

Our main interest is the unstructured P2P overlay networks, since most of the existing P2P worms target the unstructured overlay applications [2]. Naturally, due to the similar network architecture, our system can easily be deployed into the unstructured P2P overlay networks. Moreover, for structured P2P overlay networks, a subset of P2P hosts could be elected to perform the functions of Phagocytes. We aim not to change the network architecture of the structured P2P overlay networks; however, we elect Phagocytes to form an overlay to perform the defense and protection functions — this overlay acts as a security wall in a separate layer from the existing P2P overlay, thus not affecting the original P2P operations. In the next two sections, we will elaborate in detail our mechanisms of internal defense and external protection against active P2P worms.

3 Internal Defense

In this section, we first describe the active P2P worm attacks, and then, we design our internal defense mechanism.

3.1 Threat Model

Generally, active P2P worms utilize the P2P overlay to accurately retrieve the information of a few vulnerable P2P hosts, and then infect these hosts to bootstrap the worm infection. On one hand, a managed P2P host clearly knows its associated Phagocyte and its neighboring P2P hosts that are managed by the same Phagocyte; so now, an infected managed P2P host could perform the worm infection in several ways simultaneously. Firstly, the infected P2P host infects its neighboring managed P2P hosts very quickly. Secondly, the infected P2P host attempts to infect its associated Phagocyte. Lastly, the infected managed P2P host could issue P2P key queries with worms to infect many vulnerable P2P hosts managed by other Phagocytes. On the other hand, a Phagocyte could be infected as well; if so, the infected Phagocyte infects its managed P2P hosts and then its neighboring Phagocytes. As a result, in such a topological self-propagation way, the active P2P worms spread through the whole system at extraordinary speed.

3.2 Detection

Since the active P2P worms propagate based on the topological information, and do not need to probe any random IP

addresses, thus their connection failure rate should be low; moreover, the P2P worm attack traffic could easily blend into normal P2P traffic. Therefore, the active P2P worms do not exhibit easily detectable anomalies in the network traffic as normal scanning worms do.

In our system, the Phagocytes are those elected P2P hosts with the latest patches, and they can help their managed P2P hosts detect the existence of active P2P worms by monitoring these managed hosts' connection transactions and traffic volume. If a managed P2P host always sends similar queries or sets up a large number of connections, the responsible Phagocyte deduces that this managed P2P host is infected. Another pattern the Phagocytes will monitor is to determine if a portion of the managed P2P hosts have some similar behaviors such as issuing the similar queries, repeating to connect with their neighboring hosts, uploading/downloading the same files, etc., then they are considered to be infected.

Concretely, a managed P2P host's *latest* behaviors are processed into a *behavior sequence* consisting of continuous $\langle operation, payload \rangle^1$ pairs. Then, we can compute the behavior similarity between any two P2P hosts by using the *Levenshtein edit distance* [13]. Without loss of generality, we suppose that there are two behavior sequences BS_1 and BS_2 , in which $BS_{i=1,2} = \langle \langle o_{i1}, p_{i1} \rangle, \dots, \langle o_{ij}, p_{ij} \rangle, \dots, \langle o_{in}, p_{in} \rangle \rangle$, where $1 \leq j \leq n$, and n is the length of the behavior sequence. Further, we can treat each behavior sequence BS_i as the combination of the *operation sequence* $o_i = \langle o_{i1}, \dots, o_{ij}, \dots, o_{in} \rangle$ and the *payload sequence* $p_i = \langle p_{i1}, \dots, p_{ij}, \dots, p_{in} \rangle$. Now, we simultaneously *sort* the operation sequence o_2 and the payload sequence p_2 of the behavior sequence BS_2 to make the following similarity score $sim(BS_1, \widetilde{BS}_2)$ be maximum. To obtain the optimal solution, we could adopt the *maximum weighted bipartite matching* algorithm [27]; however, for efficiency, we use the *greedy* algorithm to obtain the approximate solution as an alternative.

$$\begin{cases} sim(BS_1, \widetilde{BS}_2) = \frac{\sum_{j=1}^n s(o_{1j}, \widetilde{o}_{2j}) \times s(p_{1j}, \widetilde{p}_{2j})}{n} \\ s(o_{1j}, \widetilde{o}_{2j}) = \begin{cases} 1 & \text{if } o_{1j} = \widetilde{o}_{2j} \\ 0 & \text{if } o_{1j} \neq \widetilde{o}_{2j} \end{cases} \\ s(p_{1j}, \widetilde{p}_{2j}) = 1 - \frac{d(p_{1j}, \widetilde{p}_{2j})}{\max(\text{length}(p_{1j}), \text{length}(\widetilde{p}_{2j}))} \end{cases} \quad (1)$$

Here, \widetilde{BS}_2 denotes the sorted BS_2 ; \widetilde{o}_{2j} and \widetilde{p}_{2j} denote the j^{th} item of the sorted o_2 and p_2 , respectively; $d(x, y)$ is the Levenshtein edit distance function. Finally, we treat the maximum $sim(BS_1, \widetilde{BS}_2)$ as the similarity score of the two behavior sequences. If the score exceeds a threshold θ_d , we consider the two P2P hosts perform similarly.

These detection operations are also performed between Phagocytes at the Phagocyte-tier because they could be infected as well though with latest patches. The infected

¹For simplicity, we use the abbreviation $\langle o, p \rangle$ hereafter.

Phagocytes could perform the worm propagation rapidly; however, we have the local isolation, alert propagation and software patching procedures in place to handle these infected Phagocytes after detected by their neighboring Phagocytes with the detection module as described above.

Note that, our detection mechanism is *not* a substitution for the existing worm detection mechanisms, e.g., the worm signature matching [5], but rather an effective P2P-tailored complement to them. Specifically, some *tricky* P2P worms may present the features of mild propagation rate, polymorphism, etc., so they may maliciously propagate in lower speed than the aggressive P2P worms; here, our software patching module (in section 3.5) and several existing schemes [7, 26] can help mitigate such tricky worm attacks.

Moreover, a few elaborate P2P worms, e.g., P2P-Worm.Win32.Hofox, have recently been reported to be able to kill the anti-virus/anti-worm programs on P2P hosts [2]; at the system level, some local countermeasures have been devised to protect defense tools from being eliminated, and the arms race will continue. In this paper, we assume that P2P worms cannot disable our detection module, and therefore, each Phagocyte can perform the normal detection operations as expected; so can the following modules.

3.3 Local Isolation

If a Phagocyte discovers that some of its managed P2P hosts are infected, the Phagocyte will cut off its connections with the infected P2P hosts, and ask these infected hosts to further cut off the links towards any other P2P hosts. Also, if a Phagocyte is detected (by its neighboring Phagocyte) as infected, the detecting Phagocyte immediately issues a message to ask the infected Phagocyte to cut off the connections towards the neighboring Phagocytes, and then to trigger the software patching module (in section 3.5) at the infected Phagocyte; after the software patching, these cut connections should be reestablished. With the local isolation module, our system has the capacity of self-organizing and self-healing. We utilize the local isolation to limit the impact of active P2P worms as quickly as possible.

3.4 Alert Propagation

If a worm event has been detected, i.e., any of the managed P2P hosts or neighboring Phagocytes are detected as infected, the Phagocyte propagates a worm alert to all its neighboring Phagocytes. Further, once a Phagocyte has received the worm alerts from more than a threshold θ_a of its neighboring Phagocytes, it also propagates the alert to all its neighboring Phagocytes that did not send out the alert. In general, we should appropriately tune θ_a to limit the impact area and improve the robustness against the malicious alert propagation generated by infected Phagocytes.

3.5 Software Patching

The analytical study in [24] implied that the effective software patching is feasible for an overlay network if combined with schemes to bound the worm infection rate. In our system, the security patches are published to the participating P2P hosts using the following two procedures:

Periodical patching: A patch distribution service provided by system maintainers periodically pushes the latest security patches to all Phagocytes through the underlying P2P overlay, and then these Phagocytes install and distribute them to all their managed P2P hosts. Note that, we can utilize the periodical patching to help mitigate the tricky P2P worms (in section 3.2) which are harder to be detected.

Urgent patching: When a Phagocyte is alerted of a P2P worm attack, it will immediately pull the latest patches from a system maintainer via the direct HTTP connection (for efficiency, not via the P2P overlay), and then install and disseminate them to all its managed P2P hosts.

Specifically, each patch must be signed by the system maintainer [28], so that each P2P host can verify the patch according to the signature. Note that, the zero-day vulnerabilities are not fictional, thus installing the latest patches cannot always guarantee the worm immunity. The attackers may utilize these vulnerabilities to perform deadly worm attacks. We can integrate our system with some other systems, e.g., Shield [26] and Vigilante [7], to defend against such attacks, which can be found in [4].

3.6 Preventing Attacks on External Hosts

As much as possible, the Phagocytes provide the containment of P2P worms in the P2P overlay networks. Further, we utilize the Phagocytes to implement the P2P traffic filtering mechanism which forbids any P2P connections from the P2P overlay to external hosts because such P2P connections are generally considered to be malicious — it is possible that the P2P worms ride on the P2P traffic to spread to the external hosts. We can safely make the assumption that P2P overlay traffic should be contained inside the P2P overlay boundary, and any leaked P2P traffic is abnormal. Therefore, once this leakiness is detected, the Phagocytes will perform the former procedures for local isolation, alert propagation and software patching.

4 External Protection

Our external protection mechanism aims to protect the P2P overlay network against the external worm attacks. We hide the P2P hosts' IP addresses to prevent external hosts from directly accessing the internal P2P resources. This service can be provided by a scalable distributed DNS system, e.g., CoDoNS [20]. Such DNS system returns the as-

Step 1 ($H \rightarrow P$): generate SI_H , store $SI_H \parallel SI_H \parallel (\text{null})$
Step 2 ($P \rightarrow H$): generate k , generate $SI_P \parallel SI_H$, SI_P , $k \parallel$ check SI_H , store SI_P
Step 3 ($H \rightarrow P$): retrieve $\langle SI_H, SI_P \rangle$, solve $\text{puzzle} \parallel SI_H, SI_P, k$, solution , $\text{request} \parallel$ check $\langle SI_H, SI_P \rangle$, check SI_P , check solution , store $\langle SI_H, SI_P \rangle$, process request

where

- SI_H : The session identifier of the external host.
- SI_P : The session identifier of the Phagocyte.
- k : The difficulty level of the puzzle.

Figure 2: Adaptive and Interaction-based Puzzle Scheme

sociated Phagocyte which manages the requested P2P host. Then, the Phagocyte is able to adopt our following proposed computational puzzle scheme to perform the function of access control over the requests issued by the requesting external host.

4.1 Adaptive & Interaction-based Puzzle

We propose a *novel* adaptive and interaction-based computational puzzle scheme at the Phagocytes to provide the access control over the possible external worms attacking the internal P2P overlay.

Since we are interested in how messages are processed as well as what messages are sent, for clarity and simplicity, we utilize the annotated Alice-and-Bob specification to describe our puzzle scheme. As shown in Figure 2, to gain access to the P2P overlay, an external host has to perform a lightweight interaction towards the associated Phagocyte to solve an adaptive computational puzzle; then, according to the authenticity of the puzzle solution, the Phagocyte can determine whether to process the request.

Step 1. The external host H first generates a 64-bit nonce N_H as its session identifier SI_H . Then, the external host stores SI_H and sends it to the Phagocyte.

Step 2. On receiving the message consisting of SI_H sent by the host H , the Phagocyte P adaptively adjusts the puzzle difficulty k based on the following two real-time statuses of the network environment.

- *Status of Phagocyte*: This status indicates the usage of the Phagocyte's resources, i.e., the ratio of consumed resources to total resources possessed by the Phagocyte. The more resources a Phagocyte has consumed, the harder puzzles the Phagocyte issues in the future.

- *Status of external host*: In order to mount attacks against P2P hosts effectively, malicious external hosts have no alternative but to perform the interactions and solve many computational puzzles. That is, the more connections an external host tries to establish, the higher the probability that this activity is malicious and worm-like. Hence, the more puzzles an external host has solved in the recent

period of time, the harder puzzles the Phagocyte issues to the very external host. Note that, since a malicious external host could simply spoof its IP address, in order to effectively utilize the status of external host, our computational puzzle scheme should have the capability of defending against IP spoofing attacks, which we will describe later.

Subsequently, the Phagocyte P simply generates a *unique* 64-bit session identifier SI_P for the external host according to the host's IP address IP_H (extracted from the IP header of the received message), the host's session identifier SI_H and the puzzle difficulty k , as follows:

$$SI_P = \text{HMAC}_{\text{secret}}(IP_H | SI_H | k) \quad (2)$$

Here, the HMAC is a keyed hash function for message authentication, and the *secret* is a 32-bit key which is *periodically* changed and only known to the Phagocyte itself. Such *secret* limits the time external hosts have for computing puzzle solutions, and it also guarantees that an external attacker usually does not have enough resources to pre-compute all possible solutions in step 3.

After the above generation process, the Phagocyte replies to the external host at IP_H with the host's session identifier SI_H , the Phagocyte's session identifier SI_P and the puzzle difficulty k . Once the external host has received this reply message, it first checks whether the received SI_H is really generated by itself. If the received SI_H is bogus, the external host simply drops the message; otherwise, the host stores the Phagocyte's session identifier SI_P immediately. Such reply and checking operations can effectively defend against IP spoofing attacks.

Step 3. The external host H retrieves the $\langle SI_H, SI_P \rangle$ pair as the global session identifier, and then tries to solve the puzzle according to the equation below:

$$h(SI_H, SI_P, X) = Y^{(k)} \quad (3)$$

Here, the h is a cryptographic hash function, the $Y^{(k)}$ is a hash value with the first k bits being equal to 0, and the X is the puzzle *solution*. Due to the features of hash function, the external host has no way to figure out the solution other than brute-force searching the solution space until a solution is found, even with the help of many other solved puzzles. The cost of solving the puzzle depends exponentially on the difficulty k , which can be effortlessly adjusted by the Phagocyte.

After the brute-force computation, the external host sends the Phagocyte a message including the global session identifier (i.e., the $\langle SI_H, SI_P \rangle$ pair), the puzzle difficulty, the puzzle solution and the actual *request*. Once the Phagocyte has received this message, it performs the following operations in turn:

- a) Check whether the session identifier $\langle SI_H, SI_P \rangle$ is really fresh based on the database of the past global session

identifiers. This operation can effectively defend against replay attacks.

b) Check whether the Phagocyte’s session identifier SI_P can be correctly generated according to equation 2. Specifically, this operation can additionally check whether the difficulty level k reported by the external host is the original k determined by the Phagocyte.

c) Check whether the puzzle solution is correct according to equation 3, which will also not incur significant overhead on the Phagocyte.

d) Store the global session identifier $\langle SI_H, SI_P \rangle$, and act as the overlay proxy to transmit the request submitted by the external host. Note that, in our scheme, the Phagocyte stores the session-specific data and processes the actual request only after it has verified the external host’s puzzle solution. That is, the Phagocyte does not commit its resources until the external host has demonstrated the sincerity.

Specifically in the above sequence of operations, if one operation succeeds, the Phagocyte continues to perform the next; otherwise, the Phagocyte cancels all the following operations, and the entire interaction ends. More details about the puzzle design rationale can be found in [4].

4.2 Comparison and Analysis

So far, several computational puzzle schemes [8, 10, 16] have been proposed. However, most of them consider only the status of resource providers, so they cannot reflect the network environment completely. Recently, an ingenious puzzle scheme, Portcullis [17], was proposed. In Portcullis, since a resource provider gives priority to requests containing puzzles with higher difficulty levels, to gain access to the requested resources, each resource requester, no matter legitimate or malicious, has to compete with each other and solve hard puzzles under attacks. This may influence legitimate requesters’ experiences significantly.

Compared with existing puzzle schemes, our adaptive and interaction-based computational puzzle scheme satisfies the fundamental properties of a good puzzle scheme [10]. It treats each external host *distinctively* by performing a lightweight interaction to flexibly adjust the puzzle difficulty according to the real-time statuses of the network environment. This guarantees that our computational puzzle scheme does not influence legitimate external hosts’ experiences significantly, and it also prevents a malicious external host from attacking P2P overlay without investing unbearable resources.

In real-world networks, hosts’ computation capabilities vary a lot, e.g., the time to solve a puzzle would be much different between a host with multiple fast CPUs and a host with just one slow CPU. To decrease the computational disparity, some other kinds of puzzles, e.g., memory-bound puzzle [3], could be complementary to our scheme. Note

that, with low probability, a Phagocyte may also be compromised by external worm attackers, then they could perform the topological worm propagation; here, our proposed internal defense mechanism could be employed to defend against such attacks.

5 Experimental Design

In our experiments, we first implement a prototype system, and then construct a massive-scale testbed to verify the properties of our prototype system.

5.1 Prototype System

Internal Defense. We implement an internal defense prototype system including all basic modules described in section 3. Here, a Phagocyte monitors each of its connected P2P hosts’ latest 100 requests. Firstly, if more than half of the managed P2P hosts perform similar behaviors, the responsible Phagocyte considers that the managed zone is being exploited by worm attackers. Secondly, if more than half of a Phagocyte’s neighboring Phagocytes perform the similar operations, the Phagocyte considers its neighboring Phagocytes are being exploited by worm attackers. In particular, the similarity is measured based on the equation 1 with a threshold θ_d of 0.5. Then, in the local isolation module, if a Phagocyte has detected worm attacks, the Phagocyte will cut off the associated links between the infection zone and the connected P2P hosts. Afterwards, in the alert propagation module, if a Phagocyte has detected any worm attacks, it will broadcast a worm alert to all its neighboring Phagocytes; further, if a Phagocyte receives more than half of its neighboring Phagocytes’ worm alerts, i.e., $\theta_a > 0.5$, the Phagocyte will also broadcast the alert to all its neighboring Phagocytes that did not send out the alert. Finally, in the software patching module, the Phagocytes acquire the patches from the closest one of the system maintainers (i.e., 100 online trusted Phagocytes in our testbed), and then distribute them to all their managed P2P hosts. We have not yet integrated the signature scheme into the software patching module of our prototype system.

Note that, in the above, we simply set the parameters used in our prototype system, and in real-world systems, the system designers should appropriately tune these parameters according to their specific requirements.

External Protection. We utilize our adaptive and interaction-based computational puzzle module to develop the external protection prototype system. In this prototype system, we use SHA1 as the cryptographic hash function. Generally, solving a puzzle with difficulty level k will force an external host to perform 2^{k-1} SHA1 computations on average. In particular, the difficulty level k varies between 0 and 26 in our system — this will cost an exter-

Table 1: Network Traces of Gnutella

	Trace 1	Trace 2	Trace 3	Trace 4	Trace 5	Trace 6
Number (abbr., #) of Phagocytes (Ultra-peers)	158,985	209,723	51,400	51,400	51,400	14,705
# of managed P2P hosts (Leaf-peers)	717,025	1,026,231	512,448	342,757	257,080	73,539
# of Phagocytes : # of managed P2P hosts	22.17%	20.44%	10.03%	15.00%	19.99%	20.00%
# of Phagocytes : # of all P2P hosts	18.15%	16.97%	9.12%	13.04%	16.66%	16.66%

nal host 0.0 second ($k = 0$) to 24.728 seconds ($k = 26$) on our POWER5 CPUs. In addition, the change cycle of the puzzle-related parameters is set to 5 minutes. Yet, we have not integrated our prototype system with the scalable distributed DNS system, and this work will be part of our future work.

5.2 Testbed Construction

We use the realistic network traces crawled from a million-node Gnutella network by the Cruiser [23] crawler. The dedicated massive scale Gnutella network is composed of two tiers including the ultra-peer tier and leaf-peer tier. For historical reasons, the ultra-peer tier consists of not only modern ultra-peers but also some *legacy-peers* that reside in the ultra-peer tier but cannot accept any leaf-peers. Specifically, in our experiments, the ultra-peers excluding legacy-peers perform the functions of Phagocytes, and the leaf-peers act as the managed P2P hosts.

Then, we adopt the widely accepted GT-ITM [31] to generate the transit-stub model consisting of 10,047 routers for the underlying hierarchical Internet topology. There are 10 transit domains at the top level with an average of 10 routers in each, and a link between each pair of these transit routers has a probability of 0.5. Each transit router has an average of 10 stub domains attached, and each stub has an average of 10 routers, with the link between each pair of stub routers having a probability of 0.1. There are two million end-hosts uniformly assigned to routers in the core by local area network (LAN) links. The delay of each LAN link is set to 5ms and the average delay of core links is 40ms.

Now, the crawled Gnutella networks can model the realistic P2P overlay, and the generated GT-ITM network can model the underlying Internet topology; thus, we deploy the crawled Gnutella networks upon the underlying Internet topology to simulate the realistic P2P network environment. We do not model queuing delay, packet losses and any cross network traffic because modeling such parameters would prevent the massive-scale network simulation.

As shown in Table 1, we list various Gnutella traces that we use in our experiments — with different node populations and/or different percentages of Phagocytes.

- *Trace 1*: Crawled by Cruiser on Sep. 27th, 2004.
- *Trace 2*: Crawled by Cruiser on Feb. 2nd, 2005.
- *Trace 3*: Based on trace 1, we remove a part of Phagocytes randomly; then, we remove the *isolated* Phagocytes,

i.e., these Phagocytes do not connect to any other Phagocytes; finally, we further remove the isolated managed P2P hosts, i.e., these managed P2P hosts do not connect to any Phagocytes.

- *Trace 4*: Based on trace 3, we remove a part of managed P2P hosts randomly.

- *Trace 5*: Based on trace 4, we further remove a part of managed P2P hosts randomly.

- *Trace 6*: Based on trace 1, we use the same method as described in the generation process of trace 3. In addition, we remove an extra part of managed P2P hosts.

6 Evaluation Results

6.1 Performance Metrics

In our experiments, we characterize the performance under various different circumstances by using three metrics:

- *Peak infection percentage of all P2P hosts*: The ratio of the maximum number of infected P2P hosts to the total number of P2P hosts. This metric indicates whether Phagocytes can effectively defend against internal attacks.

- *Blowup factor of latency*: This factor is the latency penalty between the external hosts and the P2P overlay via the Phagocytes and direct routing. This indicates the efficiency of our Phagocytes to filter the requests from external hosts to the P2P overlay.

- *Percentage of successful external attacks*: The ratio of the number of successful external attacks to the total number of external attacks. This metric indicates the effectiveness of our Phagocytes to prevent external hosts from attacking the P2P overlay.

6.2 Internal Defense Evaluation

In our prototype system, we model a percentage of Phagocytes and managed P2P hosts being initially *immune*, respectively; except these immune P2P hosts, the other hosts are *vulnerable*. Moreover, there are a percentage of P2P hosts being initially *infected*, which are distributed among these vulnerable Phagocytes and vulnerable managed P2P hosts uniformly at random. All the infected P2P hosts perform the active P2P worm attacks (described in section 3.1), and meanwhile, our internal defense modules deployed at each participant try to defeat such attacks. With

Table 2: Experimental Parameters (Internal Defense)

	Experiment 1	Experiment 2	Experiment 3	Experiment 4
Initial percentage (abbr., %) of immune Phagocytes	100% to 50%	95%	95%	95%
Initial % of immune managed P2P hosts	10%	0% to 30%	10%	10%
Initial infection % of all vulnerable P2P hosts	$10^{-3}\%$ to 50%	$10^{-3}\%$ to 50%	$10^{-3}\%$ to 50%	$10^{-3}\%$ to 50%
Used traces	1	1	1, 2, 5, 6	3, 4, 5







Figure 3: Impact of Immune Phagocytes

Figure 4: Impact of Immune Managed P2P Hosts

Figure 5: Impact of Network Scale

Figure 6: Impact of the Percentage of Phagocytes

different experimental parameters described in Table 2, we conduct four different experiments to evaluate the internal defense mechanism.

Experiment 1 — Impact of immune Phagocytes: With seven different initial percentages of immune Phagocytes, we fix the initial percentage of immune managed P2P hosts to 10%, and vary the number of initial infected P2P hosts so that these infected hosts make up between $10^{-3}\%$ and 50% of all the vulnerable P2P hosts. Now, we can investigate the impact of immune Phagocytes by calculating the peak infection percentage of all P2P hosts. The experimental result shown in Figure 3 demonstrates that when the initial infection percentage of all vulnerable P2P hosts is low (e.g., $< 1\%$), the Phagocytes can provide a good containment of active P2P worms; otherwise, the worm propagation is very fast, but the Phagocytes could still provide the sufficient containment — this property is also held in the following experiments. Interestingly, the initial percentage of immune Phagocytes does not influence the performance of our system significantly, i.e., the percentage of Phagocytes being initially immune has no obvious effect. This is a good property because we do not actually need to have high initial percentage of immune Phagocytes. Also, this phenomenon implies that increasing the number of immune Phagocytes does not further provide much significant defense. Thus, we can clearly conclude that the Phagocytes are effective and scalable in performing detection, local isolation, alert propagation and software patching.

Experiment 2 — Impact of immune managed P2P hosts: In this experiment, for 95% of Phagocytes being initially immune, we investigate the performance of our system with various initial percentages of immune managed P2P hosts in steps of 10%. The result shown in Figure 4 is within our expectation. The peak infection percentage of all

P2P hosts decreases with the growth of the initial percentage of immune managed P2P hosts. Actually, in real-world overlay networks, even a powerful attacker could initially control tens of thousands of overlay hosts ($1\% - 5\%$ in the X-axis); hence, we conclude that our Phagocytes have the capacity of defending against active P2P worms effectively even in a highly malicious environment.

Experiment 3 — Impact of network scale: Figure 5 plots the performance of our system in terms of different network scales. In traces 1, 2, 5 and 6, there are different node populations, but the ratios of the number of Phagocytes to the number of all P2P hosts are all around 17%. The experimental result indicates that our system can indeed help defend against active P2P worms in various overlay networks with different network scales. Furthermore, although the Phagocytes perform more effectively in smaller overlay networks (e.g., traces 5 and 6), they can still work quite well in massive-scale overlay networks with million-node participants (e.g., traces 1 and 2).

Experiment 4 — Impact of the percentage of Phagocytes: In our system, the Phagocytes perform the functions of defending against P2P worms. In this experiment, we evaluate the system performance with different percentages of Phagocytes but the same number of Phagocytes. The result in Figure 6 indicates that the higher percentage of Phagocytes, the better security defense against active P2P worms. That is, as the percentage of Phagocytes increases, we can persistently improve the security capability of defending against active P2P worms in the overlay network. Further, the experimental result also implies that we do not need to have a large number of Phagocytes to perform the defense functions — around 10% of the node population functioning as Phagocytes is sufficient for our system to provide the effective worm containment.




Figure 7: Cumulative distribution of the latency penalty between external hosts and P2P overlay via the Phagocytes and direct routing.




Figure 8: Cumulative distribution of the absolute latency difference between external hosts and P2P overlay via the Phagocytes and direct routing.




Figure 9: Effectiveness of protection against external worms attacking the P2P overlay.

6.3 External Protection Evaluation

In this section, we conduct two more experiments in our prototype system to evaluate the performance of the external protection mechanism.

Experiment 5 — Efficiency: In this experiment, we show the efficiency in terms of the latency penalty between the external hosts and the P2P overlay via the Phagocytes and direct routing. Based on trace 1, we have 100 external hosts connect to every P2P host via the Phagocytes and direct routing in turn. Then, we measure the latencies for both cases.

Figure 7 plots the measurement result of latency penalty. We can see that, if routing via the Phagocytes, about 30% and 80% of the connections between the external hosts and P2P hosts have the blowup factor of latency be less than 2 and 2.5, respectively. Figure 8 shows the corresponding absolute latency difference, from which we can further deduce that the average latency growth of more than half of these connections (via the Phagocytes) is less than 150ms. Actually, due to the interaction required by our proposed computational puzzle scheme, we would expect some latency penalty incurred by routing via the Phagocytes. With the puzzle scheme, our system can protect against external attacks effectively which we will illustrate in the next experiment. Hence, there would be a tradeoff between the efficiency and effectiveness.

Experiment 6 — Effectiveness: In this experiment, based on trace 1, we have 100 external worm attackers flood all Phagocytes in the P2P overlay. Then, we evaluate the percentage of successful external attacks to show the effectiveness of our protection mechanism against external hosts attacking the P2P overlay. For other numbers of external worm attackers, we obtain the similar experimental results.

In Figure 9, the X-axis is the attack frequency in terms of the speed of external hosts mounting worm attacks to the P2P overlay, and the Y-axis is the percentage of successful external attacks. The result clearly illustrates the effectiveness of Phagocytes in protecting the P2P overlay

from external worm attacks. Our adaptive and interaction-based computational puzzle module at the Phagocytes plays an important role in contributing to this observation. Even in an extremely malicious environment, our system is still effective. That is, to launch worm attacks, the external attackers have no alternative but to solve hard computational puzzles which will incur heavy burden on these attackers. From the Figure 9, we can also find that when the attack frequency decreases, the percentage of successful external attacks increases gradually. However, with a low attack frequency, the attackers cannot perform practical attacks. Even if a part of external attacks are mounted successfully, our internal defense mechanism can mitigate them effectively.

7 Related Work

P2P worms could exploit the perversive P2P overlays to achieve fast worm propagation, and recently, many P2P worms have already been reported to employ real-world P2P systems as their spreading platforms [2, 19, 21]. The very first work in [32] highlighted the dangers posed by P2P worms and studied the feasibility of self-defense and containment inside the P2P overlay. Afterwards, several studies [6, 19] developed mathematical models to understand the spreading behaviors of P2P worms, and showed that P2P worms, especially the active P2P worms, indeed pose more deadly threats than normal scanning worms.

Recognizing such threats, many researchers started to study the corresponding defense mechanisms. Specifically, Yu *et al.* in [30] presented a region-based active immunization defense strategy to defend against active P2P worm attacks; Freitas *et al.* in [9] utilized the diversity of participating hosts to design a worm-resistant P2P overlay, Verme, for containing possible P2P worms; moreover, in [29], Xie and Zhu proposed a partition-based scheme to proactively block the possible worm spreading as well as a connected dominating set based scheme to achieve fast patch distribution in a race with the worm, and in [28], Xie *et al.* further designed a P2P patching system through file-sharing mechanisms to

internally disseminate security patches. However, existing defense mechanisms generally focused on the internal P2P worm defense without the consideration of external worm attacks, so that they cannot provide a total worm protection for the P2P overlay systems.

8 Conclusion

In this paper, we have addressed the deadly threats posed by active P2P worms which exploit the pervasive and popular P2P applications for rapid topological worm infection. We build an immunity system that responds to the active P2P worm infection by using *Phagocytes*. The *Phagocytes* are a small subset of specially elected P2P hosts that have high immunity and can “eat” active P2P worms in the P2P overlay networks. Each *Phagocyte* manages a group of P2P hosts by monitoring their connection patterns and traffic volume. If any worm events are detected, the *Phagocyte* will invoke the internal defense strategies for local isolation, alert propagation and software patching. Besides, the *Phagocytes* provide the access control and filtering mechanisms for the communication establishment between the P2P overlay and external hosts. The *Phagocytes* forbid the P2P traffic to leak from the P2P overlay to external hosts, and further adopt a novel adaptive and interaction-based computational puzzle scheme to prevent external hosts from attacking the P2P overlay. To sum up, our holistic immunity system utilizes the *Phagocytes* to achieve both internal defense and external protection against active P2P worms. We implement a prototype system and validate its effectiveness and efficiency in massive-scale P2P overlay networks with realistic P2P network traces.

References

- [1] Gnutella. <http://www.gnutella.com/>.
- [2] Peer-to-peer worms. <http://www.viruslist.com/en/virusesdescribed?chapter=153311928>.
- [3] M. Abadi, M. Burrows, M. S. Manasse, and T. Wobber. Moderately hard, memory-bound functions. *ACM Trans. Internet Techn.*, 5(2):299–327, 2005.
- [4] Anonymized. *Phagocytes: A holistic defense and protection against active p2p worms* (tech. rep., available upon request).
- [5] D. Brumley, J. Newsome, D. X. Song, H. Wang, and S. Jha. Towards automatic generation of vulnerability-based signatures. In *IEEE Symposium on Security and Privacy*, 2006.
- [6] G. Chen and R. S. Gray. Simulating non-scanning worms on peer-to-peer networks. In *Infoscale*, 2006.
- [7] M. Costa, J. Crowcroft, M. Castro, A. I. T. Rowstron, L. Zhou, L. Zhang, and P. Barham. Vigilante: end-to-end containment of internet worms. In *SOSP*, 2005.
- [8] C. Dwork and M. Naor. Pricing via processing or combatting junk mail. In *CRYPTO*, pages 139–147, 1992.
- [9] F. Freitas, R. Rodrigues, C. Ribeiro, P. Ferreira, and L. Rodrigues. Verme: Worm containment in peer-to-peer overlays. In *IPTPS*, 2007.
- [10] A. Juels and J. G. Brainard. Client puzzles: A cryptographic countermeasure against connection depletion attacks. In *NDSS*, 1999.
- [11] J. Jung, V. Paxson, A. W. Berger, and H. Balakrishnan. Fast portscan detection using sequential hypothesis testing. In *IEEE Symposium on Security and Privacy*, 2004.
- [12] R. Kumar, D. D. Yao, A. Bagchi, K. W. Ross, and D. Rubenstein. Fluid modeling of pollution proliferation in p2p networks. In *SIGMETRICS/Performance*, pages 335–346, 2006.
- [13] V. I. Levenshtein. Binary codes capable of correcting deletions, insertions, and reversals. *Soviet Physics Doklady*, 10(8):707–710, 1966.
- [14] P. Li, M. Salour, and X. Su. A survey of internet worm detection and containment. *IEEE Communications Surveys and Tutorials*, 10(1):20–35, 2008.
- [15] E. K. Lua, J. Crowcroft, M. Pias, R. Sharma, and S. Lim. A survey and comparison of peer-to-peer overlay network schemes. *IEEE Communications Surveys and Tutorials*, 7(2):72–93, 2005.
- [16] R. C. Merkle. Secure communications over insecure channels. *Commun. ACM*, 21(4):294–299, 1978.
- [17] B. Parno, D. Wendlandt, E. Shi, A. Perrig, B. M. Maggs, and Y.-C. Hu. Portcullis: protecting connection setup from denial-of-capability attacks. In *SIGCOMM*, 2007.
- [18] N. Provos. A virtual honeypot framework. In *USENIX Security Symposium*, pages 1–14, 2004.
- [19] K. K. Ramachandran and B. Sikdar. Modeling malware propagation in gnutella type peer-to-peer networks. In *IPDPS*, 2006.
- [20] V. Ramasubramanian and E. G. Sirer. The design and implementation of a next generation name service for the internet. In *SIGCOMM*, pages 331–342, 2004.
- [21] S. Shin, J. Jung, and H. Balakrishnan. Malware prevalence in the kazaa file-sharing network. In *IMC*, 2006.
- [22] I. Stoica, R. Morris, D. R. Karger, M. F. Kaashoek, and H. Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. In *SIGCOMM*, 2001.
- [23] D. Stutzbach, R. Rejaie, and S. Sen. Characterizing unstructured overlay topologies in modern p2p file-sharing systems. In *IMC*, pages 49–62, 2005.
- [24] M. Vojnovic and A. J. Ganesh. On the effectiveness of automatic patching. In *WORM*, pages 41–50, 2005.
- [25] K. Walsh and E. G. Sirer. Experience with an object reputation system for peer-to-peer filesharing. In *NSDI*, 2006.
- [26] H. J. Wang, C. Guo, D. R. Simon, and A. Zugenmaier. Shield: vulnerability-driven network filters for preventing known vulnerability exploits. In *SIGCOMM*, 2004.
- [27] D. West. *Introduction to Graph Theory*. Prentice Hall, 2000.
- [28] L. Xie, H. Song, and S. Zhu. On the effectiveness of internal patching against file-sharing worms. In *ACNS*, 2008.
- [29] L. Xie and S. Zhu. A feasibility study on defending against ultra-fast topological worms. In *Peer-to-Peer Computing*, pages 61–70, 2007.
- [30] W. Yu, S. Chellappan, X. Wang, and D. Xuan. On defending peer-to-peer system-based active worm attacks. In *GLOBECOM*, pages 1757–1761, 2005.
- [31] E. W. Zegura, K. L. Calvert, and S. Bhattacharjee. How to model an internetwork. In *INFOCOM*, pages 594–602, 1996.
- [32] L. Zhou, L. Zhang, F. McSherry, N. Immerlica, M. Costa, and S. Chien. A first look at peer-to-peer worms: Threats and defenses. In *IPTPS*, pages 24–35, 2005.